

## WHAT IS CLAIMED IS:

Sub All  
5 1. A method for analyzing a potential cause of a change in a service, wherein service quality of the service is monitored, usage of the service is measured, and service events are detected, the method comprising:

determining a service change time window based at least in part upon a change in service quality between a first working state and a second, non-working state, and upon a change in service usage amount, the service change time window encompassing at least part of a service outage;

10 retrieving data representing a detected event and a time in which the event occurred; and

computing a probability that the detected event caused the service change based at least in part on a correlation between the event time and the service change time window.

2. The method of claim 1, wherein determining the service change time window comprises determining a service failure time window based upon the change in service quality  
15 and narrowing the service failure time window to the service change time window based upon the service usage amount measured during that service failure time window.

3. The method of claim 2, wherein the service quality is monitored through periodic polling of the service quality, and comprising determining the service failure time window as bounded by a polled point of the first working state and a polled point of the second,  
20 non-working state.

4. The method of claim 1, wherein computing the probability comprises computing the probability using at least in part a time weighting function which decreases exponentially with the distance between the event time and the service change time window.

5. The method of claim 1, comprising determining whether one or more other events of a type identical to the detected event occurred, and wherein computing the probability comprises computing the probability using at least in part a false occurrence weighting function which decreases the probability of the detected event as the cause of the service change for instances in which the detected event occurred outside the service change time window.

6. The method of claim 1, comprising storing historical data associating occurrences of prior events with prior service changes, and wherein computing the probability that the detected event caused the service change comprises computing the probability based at least in part on the historical data.

7. The method of claim 6, wherein storing historical data comprises storing data representing instances in which prior events occurred within prior service change time windows, and wherein computing the probability that the detected event caused the service change comprises using at least in part a positive occurrence weighting function which increases the probability of the detected event as the cause of the service change based on instances in the historical data in which a prior event of a type identical to the detected event occurred within a prior service change time window.

8. The method of claim 6, wherein storing historical data comprises storing data representing instances in which prior events were identified as having caused prior service changes, and wherein computing the probability that the detected event caused the service change comprises using at least in part a historical weighting function which increases the probability of the detected event as the cause of the service change based on instances in the historical data in which a prior event of a type identical to the detected event was identified as having caused a prior service change.

9. The method of claim 1, comprising retrieving data representing a plurality of detected events and corresponding event times, and wherein computing the probability comprises computing probabilities for each of the plurality of detected events.

10. The method of claim 9, wherein computing probabilities comprises  
5 computing the probabilities such that the total of all computed probabilities is 1.

11. The method of claim 1, wherein the service comprises service over a communication network and wherein the detected event comprises a network event.

12. The method of claim 1, wherein the service comprises service provided by an application program and wherein the detected event comprises an application program event.

10 13. The method of claim 1, wherein the service change is a service outage,  
comprising determining the service change time window as a change in service quality from the first working state to the second, non-working state.

14. The method of claim 1, wherein the service change is a service recovery,  
comprising determining the service change time window as a change in service quality from the  
15 second, non-working state to the first, working state.

15. The method of claim 1, wherein determining the service change time window comprises detecting a change in service quality by detecting a step change in measured usage.

16. A method for analyzing potential causes of a service change, the method comprising:

20 determining a service change time window encompassing a change of service between a first working state and a service outage, the service change being determined at least in part based on measured service usage levels;

detecting occurrences of a set of events within a given time prior to and during the service change time window, each occurrence of an event being associated with a time at which the event occurred; and

computing a probability distribution for the set of events, which probability distribution determines for each event in the set the probability that the detected event caused the service change, the probability distribution being based at least in part on relations between the time of each event occurrence and the service change time window.

17. The method of claim 16, wherein computing the probability distribution for the set of events comprises computing the probability distribution using a first weighting function which is the product of two or more second weighting functions.

18. The method of claim 16, wherein the two or more second functions are selected from the group consisting of:

a time weighting function which decreases exponentially the probability of a given event as the cause of the service change with the distance between the given event time and the service change time window;

a false occurrence weighting function which decreases the probability of a given event as the cause of the service change for instances in which events of the same type as the given event occurred outside the service change time window;

a positive occurrence weighting function which increases the probability of a given event as the cause of the service change based on instances stored in a historical database in which events of the same type as the given event occurred within a prior service change time window; and

a historical weighting function which increases the probability of a given event as the cause of the service change based on instances in the historical database in which events of the same type as the given event were identified as having caused a prior service outage.

19. The method of claim 18, wherein the step of computing the probability  
5 distribution comprises using a first weighting function which is the product of the time weighting function, false occurrence weighting function, positive occurrence weighting function, and user weighting function.

20. The method of claim 16, comprising monitoring service quality, and wherein  
10 determining the service change time window comprises determining a service failure time window based upon a change in monitored service quality and narrowing the service failure time window to the service change time window based upon the service usage amount measured during that service failure time window.

21. The method of claim 20, wherein the service quality is monitored through  
15 periodic polling of the service quality, and comprising determining the service failure time window as bounded by a polled point of the first working state and a polled point of the second, non-working state.

22. The method of claim 16, comprising computing the probability distribution  
such that the total of all probabilities in the distribution is 1.

23. The method of claim 16, wherein the service comprises service over a  
20 communication network and wherein the detected events comprise network events.

24. The method of claim 16, wherein the service comprises service provided by  
an application program and wherein the detected events comprise application program events.

25. The method of claim 16, wherein the service change is a service outage, comprising determining the service change time window as a change in service from the first working state to the second, non-working state.

26. The method of claim 16, wherein the service change is a service recovery, comprising determining the service change time window as a change in service from the second, non-working state to the first, working state.

27. The method of claim 1, wherein determining the service change time window comprises detecting a step change in measured usage.

28. A network monitoring system comprising:

a service monitor for monitoring quality of service on the network;

a usage meter for measuring usage of the network;

an event detector for detecting network events and times at which the network events occur; and

a probable cause engine, coupled to receive data from the service monitor, usage meter, and the event detector, for:

setting a service change time window based upon data received from the service monitor or usage meter, the service change time window encompassing at least part of an occurrence of a service outage in the network; and

determining which of the network events detected by the event detector is the most likely cause of a service change based at least in part of the relations of the detected network event times to the service change time window.

29. A computer readable medium storing program code for, when executed, causing a computer to perform a method for analyzing a potential cause of an change in a

service, wherein service quality of the service is monitored, usage amount of the service is measured, and service events are detected, the method comprising:

determining a service change time window based at least in part upon a change in service quality between a first working state and a second, non-working state, and upon a change in service usage amount, the service change time window encompassing at least part of a service outage;

retrieving data representing a detected event and a time in which the event occurred; and

computing a probability that the detected event caused the service change based at least in part on a correlation between the event time and the service change time window.

30. A method for quantifying the effect of an outage in a service over a first period of time, the method comprising:

measuring usage of the service over time;

defining a cost of outage time window comprising the first time period and a second time period following the first time period; and

computing a cost of outage as the difference between the measured service usage during the cost of outage time window with service usage measured during a comparison window, the comparison window being substantially equal in time to that of the cost of outage time window and reflecting a similar period of service activity as that of the cost of outage time window without having a service outage.

31. The method of claim 30, comprising determining the second period of time to be a time in which the measured service usage returns to within a given percentage of a normal service usage.

32. The method of claim 30, comprising determining the second period of time to be the shorter of (1) a time in which the measured service usage returns to within a given percentage of a normal service usage and (2) a maximum time period.

33. The method of claim 30, wherein computing the cost of outage comprises  
5 computing the difference in units of service usage.

34. The method of claim 33, wherein the service is a communication service conveying a plurality of messages, the method comprising computing the cost of outage in numbers of messages conveyed.

35. The method of claim 33, wherein the service is a network server providing  
10 data items in response to requests therefor, the method comprising computing the cost of outage in numbers of requests received or data items provided by a server on the network.

36. The method of claim 33, comprising converting the computed units of cost of service outage to a monetary value.

37. The method of claim 36, wherein converting the computed units of cost of  
15 service outage comprises multiplying the units of cost of service outage by a first monetary value per unit of usage.

38. The method of claim 30, comprising comparing the cost of outage to a second cost of outage value for a different service and prioritizing the outages based on the compared costs.

39. The method of claim 30, comprising computing the difference between the  
20 monitored service usage following the cost of outage time window and a normal service usage level to thereby measure a long term effect of the service outage.



40. A method for quantifying the effect of an outage in a service, the method comprising:

measuring usage amounts of the service during a period of the service outage and a second period following the service outage;

5 comparing the measured usage amounts to normal usage amounts measured under similar service conditions for a similar period of time where no service outage occurs; and determining a level of loss of service due to the service outage based on the comparison.

41. The method of claim 40, comprising defining the second period as the shorter  
10 of a time period in which measured service usage amounts return to within a given range of normal usage amounts and a predefined maximum time period.

42. The method of claim 40, wherein measuring service usage amounts comprises measuring service usage amounts in terms of units of service usage.

43. The method of claim 42, wherein the service is a communication service  
15 conveying a plurality of messages, comprising measuring service usage amounts in terms of number of messages conveyed by the system.

44. The method of claim 42, wherein the service is a network server providing data items in response to requests therefor, comprising measuring service usage amounts in terms of numbers of requests received or data items provided by a server on the network.

20 45. The method of claim 40, wherein determining the level of service loss comprises determining that substantially no loss of service occurred due to the outage based on the measured service usage amounts and normal service usage amounts being substantially equal.

46. The method of claim 40, comprising:

measuring service usage amounts during a third period following the second period;

comparing the measured third period service usage amounts to normal usage amounts measured under similar service conditions for a similar period of time; and

5 determining a long term effect on the service due to the service outage based on the comparison.

47. A computer readable medium storing program code which, when executed, causes a computer to perform a method for quantifying the effect of an outage in a service over a first period of time, the method comprising:

10 measuring service usage over time;

defining a cost of outage time window comprising the first time period and a second time period; and

computing a cost of outage as the difference between the measured level of service usage during the cost of outage time window with a level of usage in a comparison window, the comparison window being substantially equal in time to the cost of outage time window and reflecting a similar period of service activity as the cost of outage time window without having a service outage.

15 48. A method for predicting a cost of an outage of a service, the method comprising:

20 measuring time duration for and service usage during the outage;

comparing the measured usage amounts to normal usage amounts measured under similar service conditions for a similar period of time where no service outage occurs, to thereby determine a usage loss amount; and

computing a predicted cost of the outage based at least upon a cost component,  
the cost component comprising a function of the measured time of the outage and measured  
usage loss amount.

49. The method of claim 48, comprising measuring service usage on an ongoing  
5 basis and detecting the onset of the service outage using the measured service usage.

50. The method of claim 49, wherein detecting the onset of the service outage  
comprises detecting a step change in service usage.

51. The method of claim 48, comprising monitoring quality of the service and  
detecting the onset of a service outage based upon the service quality.

52. The method of claim 51, wherein monitoring service quality comprises  
10 monitoring service quality through periodic polling of the service quality, and wherein detecting  
the onset of a service outage comprises detecting the outage onset as bounded by a polled point  
of a first, working state and a polled point of a second, non-working state.

53. The method of claim 48, wherein computing the predicted cost of the outage  
15 comprises using a service demand cost component representing an affect on service usage based  
upon the duration of an outage.

54. The method of claim 53, wherein using the service demand cost component  
comprises multiplying the measured usage loss by a usage loss curve which is a function of time  
duration of an outage and represents a predicted percentage usage due to an outage based on time  
20 duration of the outage.

55. The method of claim 54, comprising generating the usage loss curve using  
historical data derived from prior service outages.

56. The method of claim 48, wherein computing the predicted cost of the outage comprises using a customer retention cost component representing a number or percentage of customers lost due to the outage.

57. The method of claim 48, wherein computing the predicted cost of the outage comprises using an agreement penalty component representing penalties arising under one or more service agreements due to a service outage.

58. The method of claim 48, wherein computing the predicted cost of the outage comprises computing the cost in units of service usage.

59. The method of claim 58, comprising converting the computed units of predicted cost to a monetary value by multiplying the units of predicted cost by a first monetary value per unit of usage.

60. The method of claim 48, comprising comparing the predicted cost of service outage to a second predicted cost of outage value for a different service and prioritizing the outages based on the compared costs.

61. A network monitoring system comprising:  
 a usage meter for measuring usage of a service on the network;  
 an event detector for detecting network events and times at which the network events occur;  
 a probable cause engine, coupled to receive data from the usage meter and the event detector, for determining which of the network events detected by the event detector is the most likely cause of a service outage based at least in part of the relations of the detected network event times to a service change time window, the service change time window encompassing at least part of an occurrence of the service outage in the network; and

